

EXHIBIT 4

FILED UNDER SEAL



US 20080025137A1

(19) **United States**(12) **Patent Application Publication****Rajan et al.**(10) **Pub. No.: US 2008/0025137 A1**(43) **Pub. Date: Jan. 31, 2008**(54) **SYSTEM AND METHOD FOR SIMULATING AN ASPECT OF A MEMORY CIRCUIT**(22) Filed: **Jul. 31, 2006****Publication Classification**

(75) Inventors: **Suresh Natarajan Rajan**, San Jose, CA (US); **Keith R. Schakel**, San Jose, CA (US); **Michael John Sebastian Smith**, Palo Alto, CA (US); **David T. Wang**, San Jose, CA (US); **Frederick Daniel Weber**, San Jose, CA (US)

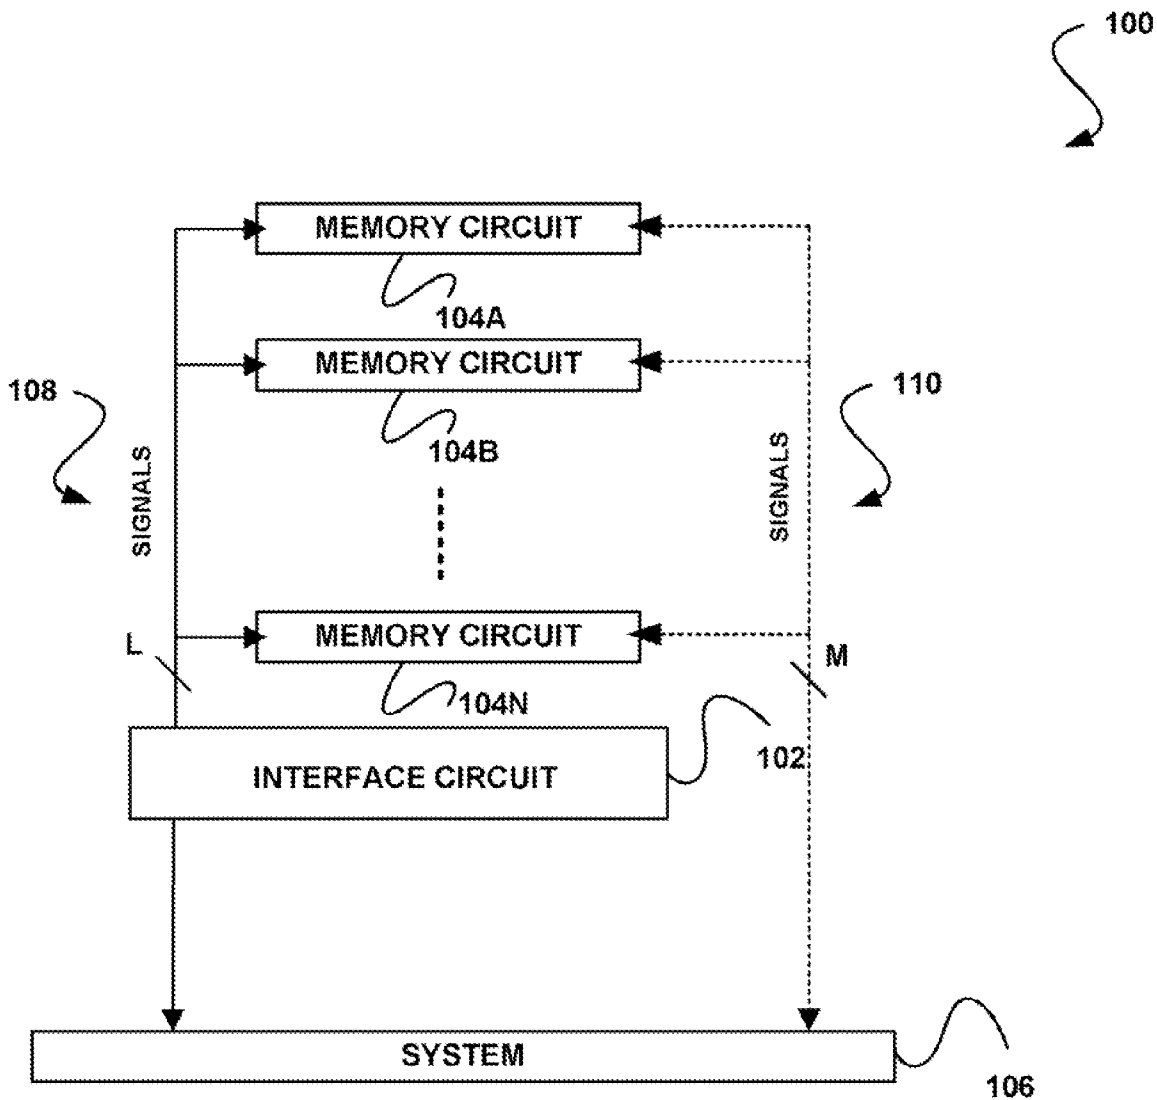
(51) **Int. Cl.**
G11C 8/00 (2006.01)

(52) **U.S. Cl.** **365/239**

(57) **ABSTRACT**

A system and method are provided for simulating an aspect of a memory circuit. Included is an interface circuit that is in communication with a plurality of memory circuits and a system. Such interface circuit is operable to interface the memory circuits and the system for simulating at least one memory circuit with at least one aspect that is different from at least one aspect of at least one of the plurality of memory circuits. In accordance with various embodiments, such aspect may include a signal, a capacity, a timing, and/or a logical interface.

Correspondence Address:

ZILKA-KOTAB, PC**P.O. BOX 721120****SAN JOSE, CA 95172-1120**(73) Assignee: **METARAM, INC.**(21) Appl. No.: **11/461,427**




FIGURE 1




FIGURE 2A




FIGURE 2B




FIGURE 2C




FIGURE 2D




FIGURE 2E

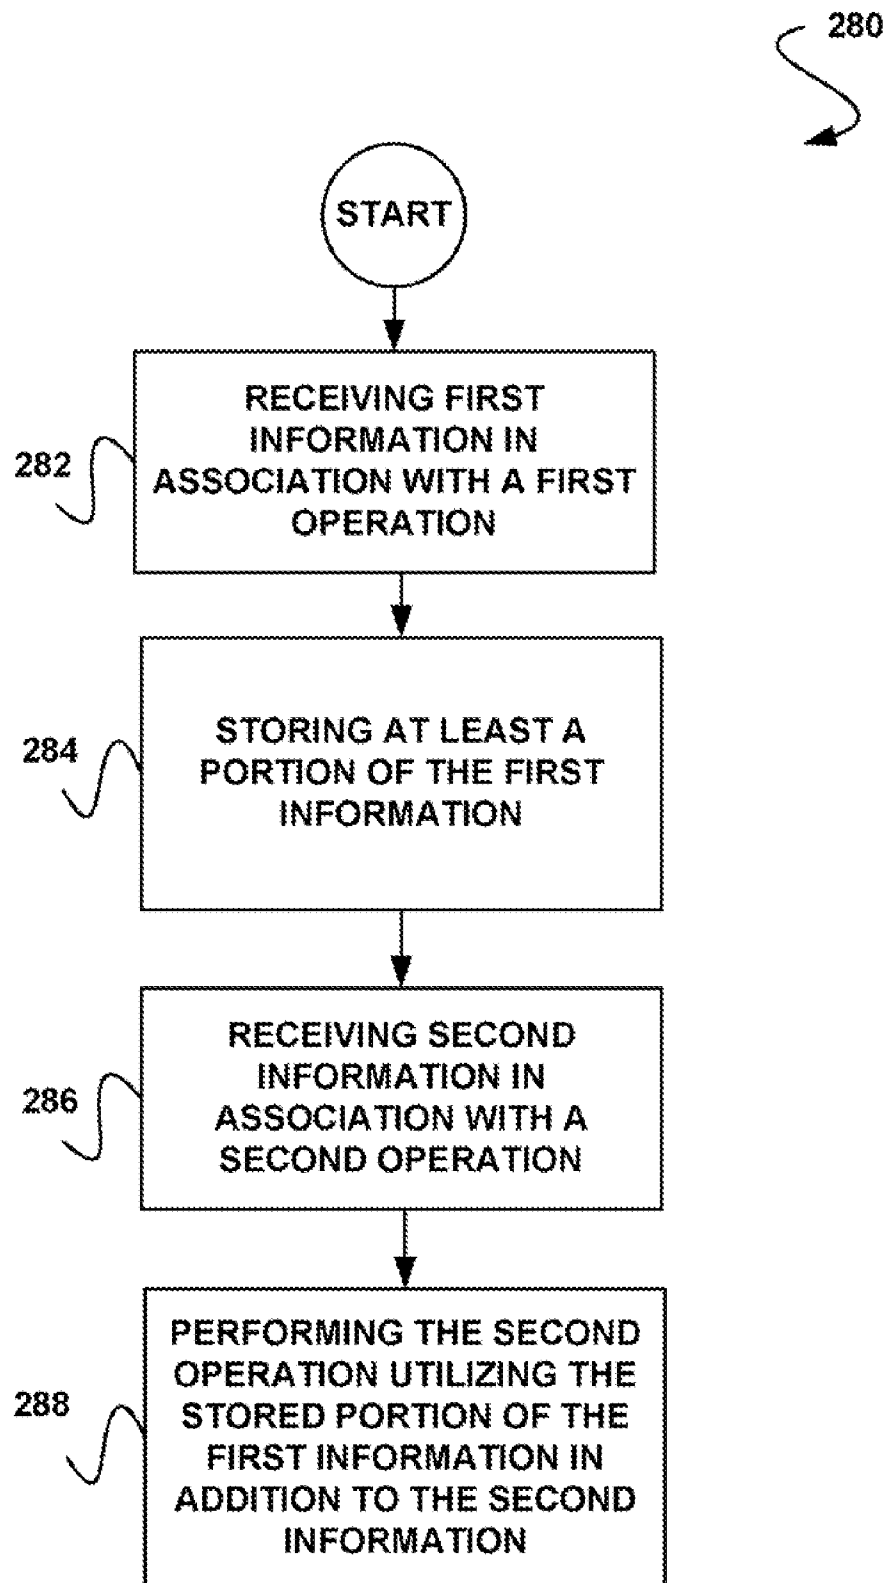


FIGURE 2F




FIGURE 3




FIGURE 4




FIGURE 5




FIGURE 6




FIGURE 7




FIGURE 8

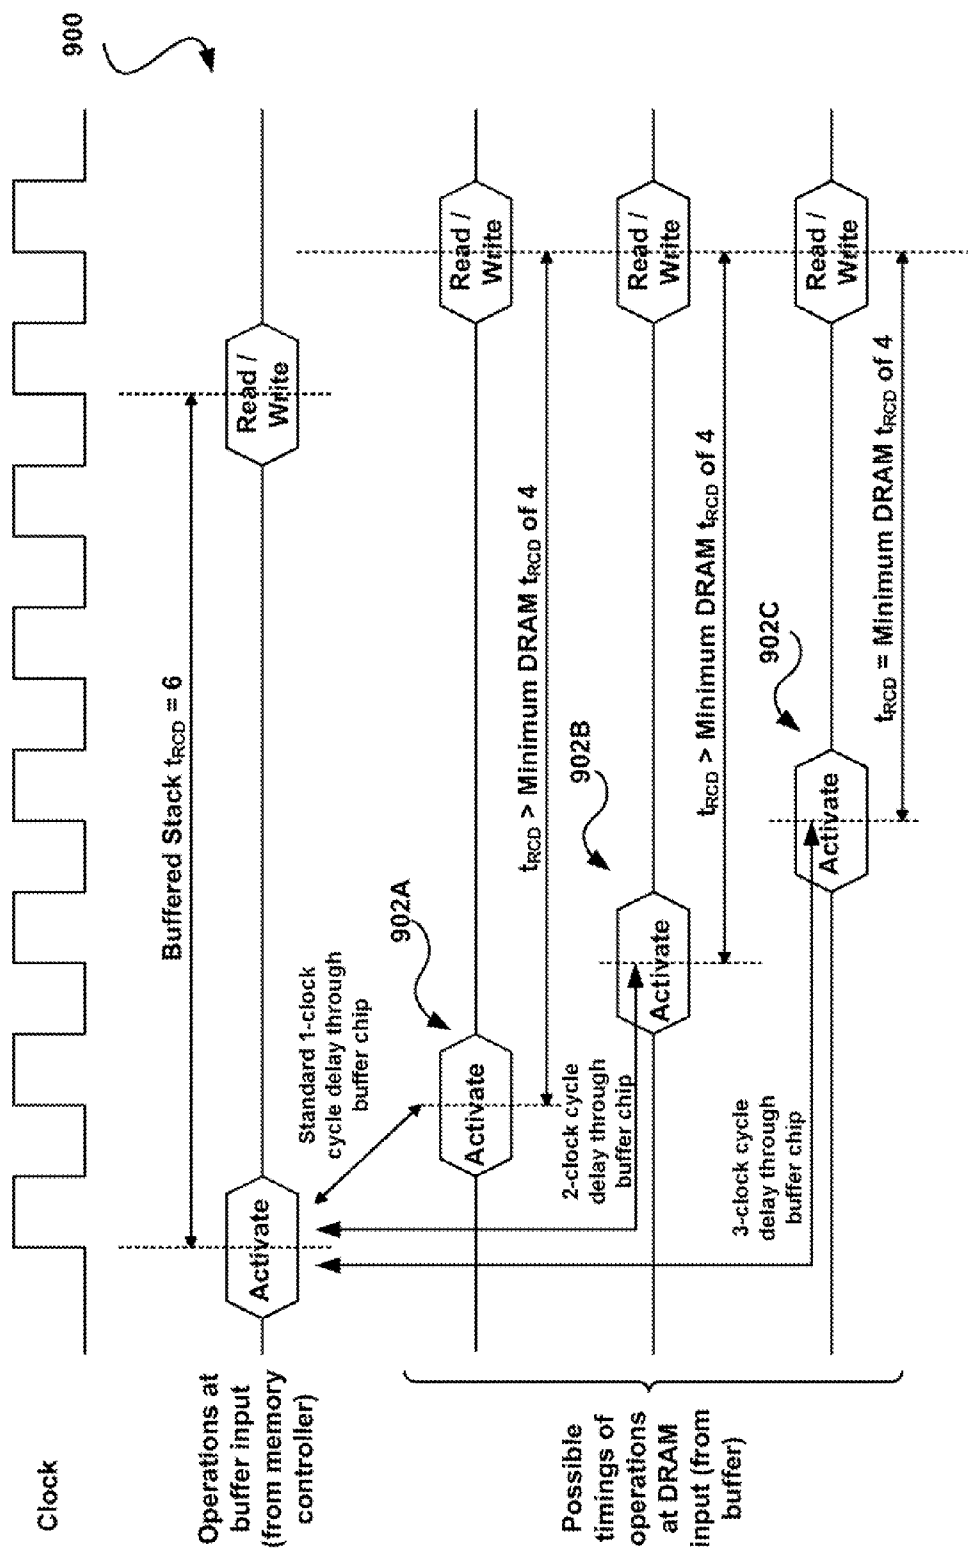


FIGURE 9A




FIGURE 9B

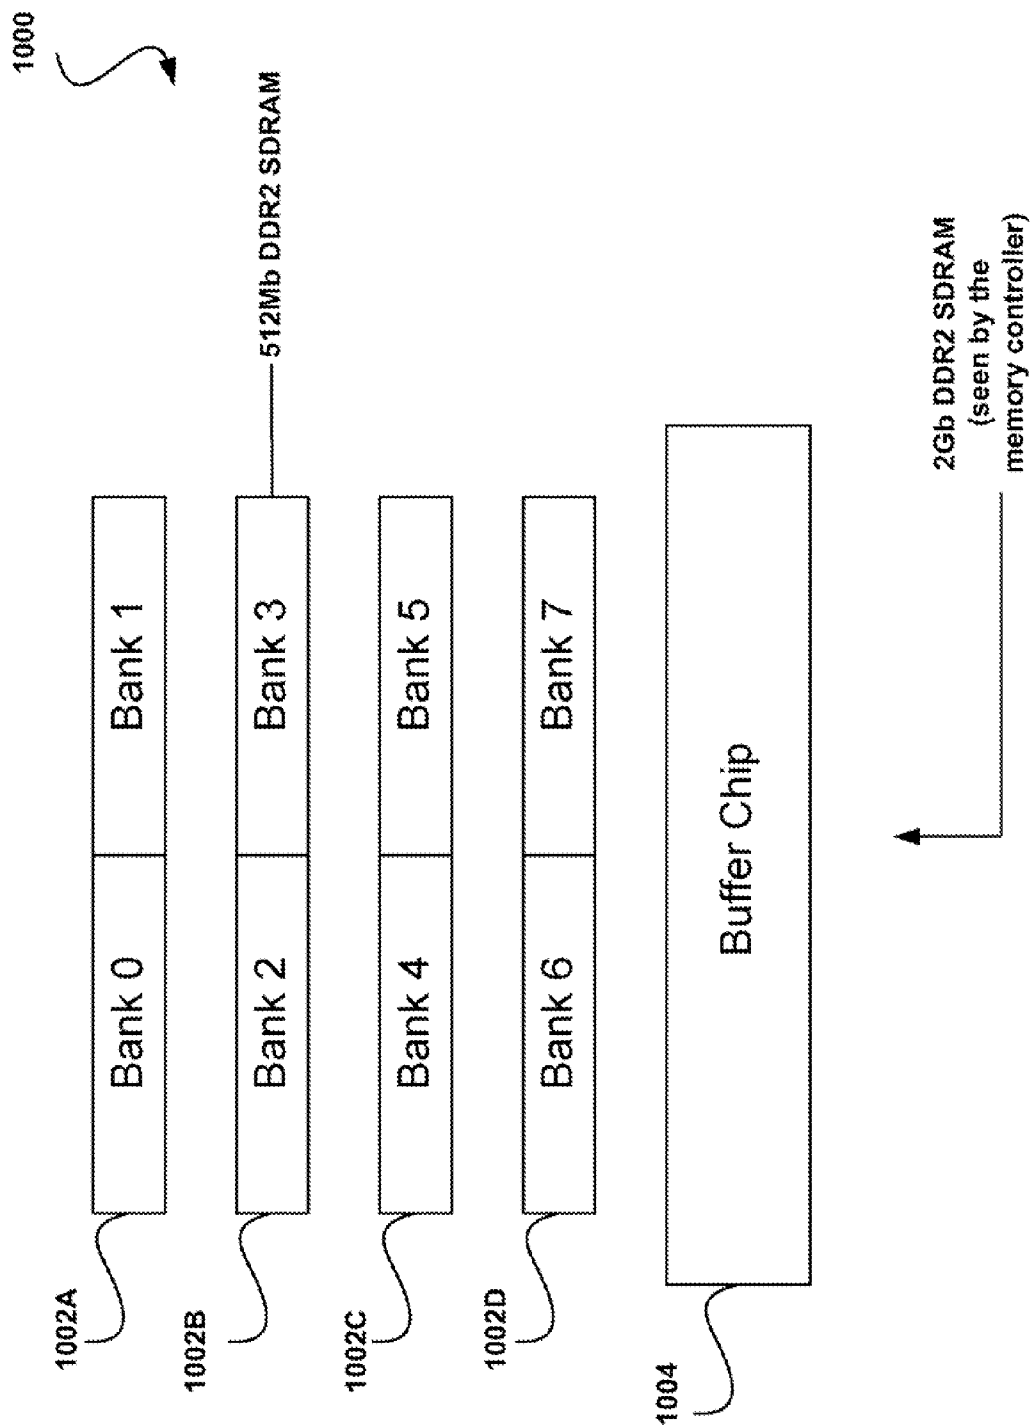


FIGURE 10

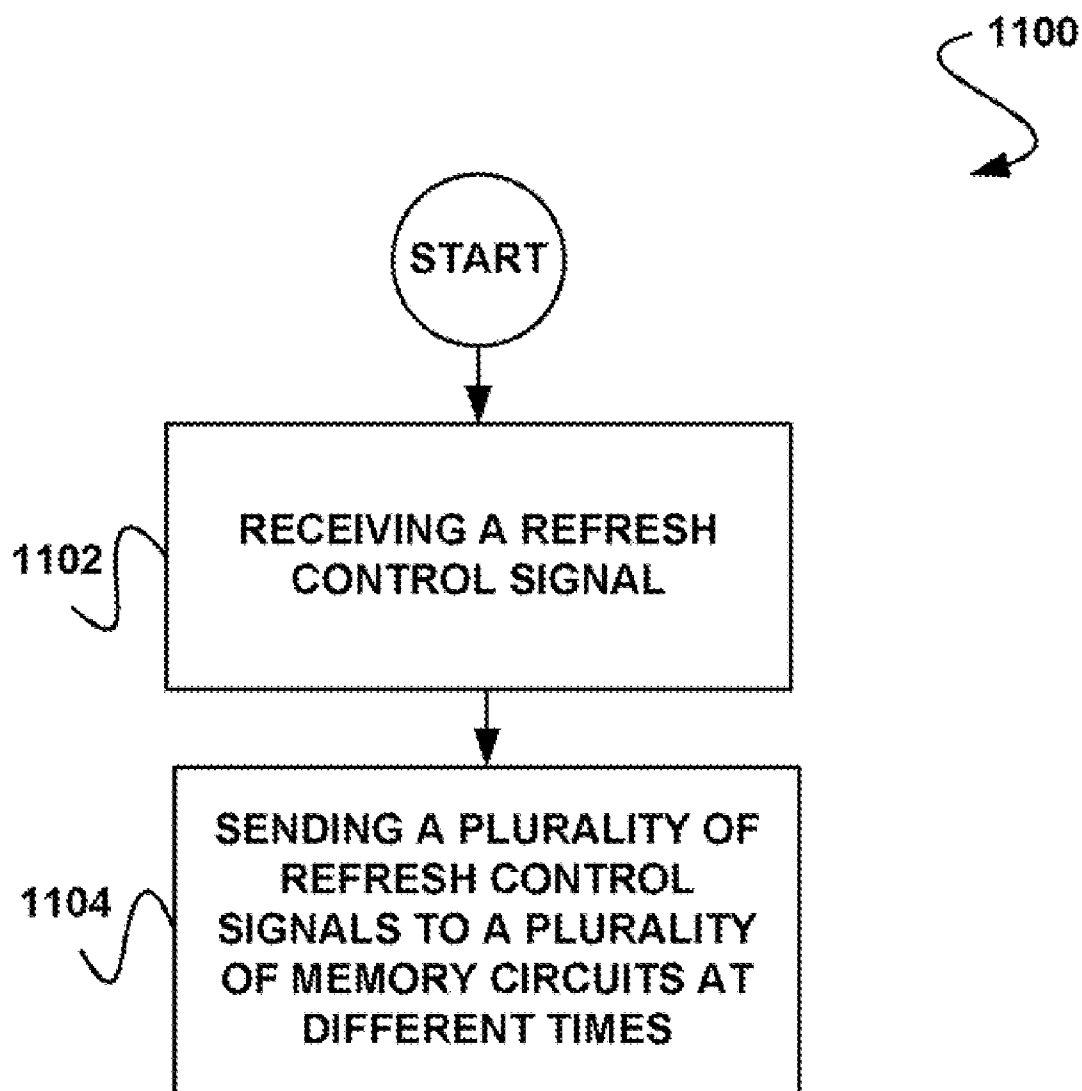


FIGURE 11

US 2008/0025137 A1

Jan. 31, 2008

1

SYSTEM AND METHOD FOR SIMULATING AN ASPECT OF A MEMORY CIRCUIT

FIELD OF THE INVENTION

[0001] The present invention relates to memory, and more particularly to multiple-memory circuit systems.

BACKGROUND

[0002] The memory capacity requirements of computers in general, and servers in particular, are increasing rapidly due to various trends such as 64-bit processors and operating systems, multi-core processors, virtualization, etc. However, other industry trends such as higher memory bus speeds and small form factor machines, etc. are reducing the number of memory module slots in such systems. Thus, a need exists in the industry for large capacity memory circuits to be used in such systems.

[0003] However, there is also an exponential relationship between a capacity of monolithic memory circuits and a price associated therewith. As a result, large capacity memory modules may be cost prohibitive. To this end, the use of multiple smaller capacity memory circuits is a cost-effective approach to increasing such memory capacity.

SUMMARY

[0004] A system and method are provided for simulating an aspect of a memory circuit. Included is an interface circuit that is in communication with a plurality of memory circuits and a system. Such interface circuit is operable to interface the memory circuits and the system for simulating at least one memory circuit with at least one aspect that is different from at least one aspect of at least one of the plurality of memory circuits. In accordance with various embodiments, such aspect may include a signal, a capacity, a timing, and/or a logical interface.

BRIEF DESCRIPTION OF THE DRAWINGS

[0005] FIG. 1 illustrates a multiple memory circuit framework, in accordance with one embodiment.

[0006] FIGS. 2A-2E show various configurations of a buffered stack of dynamic random access memory (DRAM) circuits with a buffer chip, in accordance with various embodiments.

[0007] FIG. 2F illustrates a method for storing at least a portion of information received in association with a first operation for use in performing a second operation, in accordance with still another embodiment.

[0008] FIG. 3 shows a high capacity dual in-line memory module (DIMM) using buffered stacks, in accordance with still yet another embodiment.

[0009] FIG. 4 shows a timing design of a buffer chip that makes a buffered stack of DRAM circuits mimic longer column address strobe (CAS) latency DRAM to a memory controller, in accordance with another embodiment.

[0010] FIG. 5 shows the write data timing expected by DRAM in a buffered stack, in accordance with yet another embodiment.

[0011] FIG. 6 shows write control signals delayed by a buffer chip, in accordance with still yet another embodiment.

[0012] FIG. 7 shows early write data from an advanced memory buffer (AMB), in accordance with another embodiment.

[0013] FIG. 8 shows address bus conflicts caused by delayed write operations, in accordance with yet another embodiment.

[0014] FIGS. 9A-B show variable delays of operations through a buffer chip, in accordance with another embodiment.

[0015] FIG. 10 shows a buffered stack of four 512 Mb DRAM circuits mapped to a single 2 Gb DRAM circuit, in accordance with yet another embodiment.

[0016] FIG. 11 illustrates a method for refreshing a plurality of memory circuits, in accordance with still yet another embodiment.

DETAILED DESCRIPTION

[0017] FIG. 1 illustrates a multiple memory circuit framework 100, in accordance with one embodiment. As shown, included are an interface circuit 102, a plurality of memory circuits 104A, 104B, 104N, and a system 106. In the context of the present description, such memory circuits 104A, 104B, 104N may include any circuit capable of serving as memory.

[0018] For example, in various embodiments, one or more of the memory circuits 104A, 104B, 104N may include a monolithic memory circuit. For instance, such monolithic memory circuit may take the form of dynamic random access memory (DRAM). Such DRAM may take any form including, but not limited to synchronous (SDRAM), double data rate synchronous (DDR DRAM, DDR2 DRAM, DDR3 DRAM, etc.), quad data rate (QDR DRAM), direct RAM-BUS (DRDRAM), fast page mode (FPM DRAM), video (VDRAM), extended data out (EDO DRAM), burst EDO (BEDO DRAM), multibank (MDRAM), synchronous graphics (SGRAM), and/or any other type of DRAM. Of course, one or more of the memory circuits 104A, 104B, 104N may include other types of memory such as magnetic random access memory (MRAM), intelligent random access memory (IRAM), distributed network architecture (DNA) memory, window random access memory (WRAM), flash memory (e.g. NAND, NOR, or others, etc.), pseudostatic random access memory (PSRAM), wetware memory, and/or any other type of memory circuit that meets the above definition.

[0019] In additional embodiments, the memory circuits 104A, 104B, 104N may be symmetrical or asymmetrical. For example, in one embodiment, the memory circuits 104A, 104B, 104N may be of the same type, brand, and/or size, etc. Of course, in other embodiments, one or more of the memory circuits 104A, 104B, 104N may be of a first type, brand, and/or size; while one or more other memory circuits 104A, 104B, 104N may be of a second type, brand, and/or size, etc. Just by way of example, one or more memory circuits 104A, 104B, 104N may be of a DRAM type, while one or more other memory circuits 104A, 104B, 104N may be of a flash type. While three or more memory circuits 104A, 104B, 104N are shown in FIG. 1 in accordance with one embodiment, it should be noted that any plurality of memory circuits 104A, 104B, 104N may be employed.

[0020] Strictly as an option, the memory circuits 104A, 104B, 104N may or may not be positioned on at least one dual in-line memory module (DIMM) (not shown). In various embodiments, the DIMM may include a registered DIMM (R-DIMM), a small outline-DIMM (SO-DIMM), a fully buffered-DIMM (FB-DIMM), an un-buffered DIMM,

US 2008/0025137 A1

Jan. 31, 2008

2

etc. Of course, in other embodiments, the memory circuits **104A**, **104B**, **104N** may or may not be positioned on any desired entity for packaging purposes.

[0021] Further in the context of the present description, the system **106** may include any system capable of requesting and/or initiating a process that results in an access of the memory circuits **104A**, **104B**, **104N**. As an option, the system **106** may accomplish this utilizing a memory controller (not shown), or any other desired mechanism. In one embodiment, such system **106** may include a host system in the form of a desktop computer, lap-top computer, server, workstation, a personal digital assistant (PDA) device, a mobile phone device, a television, a peripheral device (e.g. printer, etc.). Of course, such examples are set forth for illustrative purposes only, as any system meeting the above definition may be employed in the context of the present framework **100**.

[0022] Turning now to the interface circuit **102**, such interface circuit **102** may include any circuit capable of indirectly or directly communicating with the memory circuits **104A**, **104B**, **104N** and the system **106**. In various optional embodiments, the interface circuit **102** may include one or more interface circuits, a buffer chip, etc. Embodiments involving such a buffer chip will be set forth hereinafter during reference to subsequent figures. In still other embodiments, the interface circuit **102** may or may not be manufactured in monolithic form.

[0023] While the memory circuits **104A**, **104B**, **104N**, interface circuit **102**, and system **106** are shown to be separate parts, it is contemplated that any of such parts (or portions thereof) may or may not be integrated in any desired manner. In various embodiments, such optional integration may involve simply packaging such parts together (e.g. stacking the parts, etc.) and/or integrating them monolithically. Just by way of example, in various optional embodiments, one or more portions (or all, for that matter) of the interface circuit **102** may or may not be packaged with one or more of the memory circuits **104A**, **104B**, **104N** (or all, for that matter). Different optional embodiments which may be implemented in accordance with the present multiple memory circuit framework **100** will be set forth hereinafter during reference to FIGS. 2A-2E, and 3 et al.

[0024] In use, the interface circuit **102** may be capable of various functionality, in the context of different embodiments. More illustrative information will now be set forth regarding such optional functionality which may or may not be implemented in the context of such interface circuit **102**, per the desires of the user. It should be strongly noted that the following information is set forth for illustrative purposes and should not be construed as limiting in any manner. For example, any of the following features may be optionally incorporated with or without the exclusion of other features described.

[0025] For instance, in one optional embodiment, the interface circuit **102** interfaces a plurality of signals **108** that are communicated between the memory circuits **104A**, **104B**, **104N** and the system **106**. As shown, such signals may, for example, include address/control/clock signals, etc. In one aspect of the present embodiment, the interfaced signals **108** may represent all of the signals that are communicated between the memory circuits **104A**, **104B**, **104N** and the system **106**. In other aspects, at least a portion of signals **110** may travel directly between the memory circuits

104A, **104B**, **104N** and the system **106** or component thereof [e.g. register, advanced memory buffer (AMB), memory controller, or any other component thereof, where the term component is defined hereinbelow]. In various embodiments, the number of the signals **108** (vs. a number of the signals **110**, etc.) may vary such that the signals **108** are a majority or more ($L > M$), etc.

[0026] In yet another embodiment, the interface circuit **102** may be operable to interface a first number of memory circuits **104A**, **104B**, **104N** and the system **106** for simulating at least one memory circuit of a second number. In the context of the present description, the simulation may refer to any simulating, emulating, disguising, transforming, converting, and/or the like that results in at least one aspect (e.g. a number in this embodiment, etc.) of the memory circuits **104A**, **104B**, **104N** appearing different to the system **106**. In different embodiments, the simulation may be electrical in nature, logical in nature, protocol in nature, and/or performed in any other desired manner. For instance, in the context of electrical simulation, a number of pins, wires, signals, etc. may be simulated, while, in the context of logical simulation, a particular function may be simulated. In the context of protocol, a particular protocol (e.g. DDR3, etc.) may be simulated.

[0027] In still additional aspects of the present embodiment, the second number may be more or less than the first number. Still yet, in the latter case, the second number may be one, such that a single memory circuit is simulated. Different optional embodiments which may employ various aspects of the present embodiment will be set forth hereinafter during reference to FIGS. 2A-2E, and 3 et al.

[0028] In still yet another embodiment, the interface circuit **102** may be operable to interface the memory circuits **104A**, **104B**, **104N** and the system **106** for simulating at least one memory circuit with at least one aspect that is different from at least one aspect of at least one of the plurality of the memory circuits **104A**, **104B**, **104N**. In accordance with various aspects of such embodiment, such aspect may include a signal, a capacity, a timing, a logical interface, etc. Of course, such examples of aspects are set forth for illustrative purposes only and thus should not be construed as limiting, since any aspect associated with one or more of the memory circuits **104A**, **104B**, **104N** may be simulated differently in the foregoing manner.

[0029] In the case of the signal, such signal may refer to a control signal (e.g. an address signal; a signal associated with an activate operation, precharge operation, write operation, read operation, a mode register write operation, a mode register read operation, a refresh operation; etc.), a data signal, a logical or physical signal, or any other signal for that matter. For instance, a number of the aforementioned signals may be simulated to appear as fewer or more signals, or even simulated to correspond to a different type. In still other embodiments, multiple signals may be combined to simulated another signal. Even still, a length of time in which a signal is asserted may be simulated to be different.

[0030] In the case of protocol, such may, in one exemplary embodiment, refer to a particular standard protocol. For example, a number of memory circuits **104A**, **104B**, **104N** that obey a standard protocol (e.g. DDR2, etc.) may be used to simulate one or more memory circuits that obey a different protocol (e.g. DDR3, etc.). Also, a number of memory circuits **104A**, **104B**, **104N** that obey a version of protocol (e.g. DDR2 with 3-3-3 latency timing, etc.) may be

US 2008/0025137 A1

Jan. 31, 2008

3

used to simulate one or more memory circuits that obey a different version of the same protocol (e.g. DDR2 with 5-5-5 latency timing, etc.).

[0031] In the case of capacity, such may refer to a memory capacity (which may or may not be a function of a number of the memory circuits **104A**, **104B**, **104N**; see previous embodiment). For example, the interface circuit **102** may be operable for simulating at least one memory circuit with a first memory capacity that is greater than (or less than) a second memory capacity of at least one of the memory circuits **104A**, **104B**, **104N**.

[0032] In the case where the aspect is timing-related, the timing may possibly relate to a latency (e.g. time delay, etc.). In one aspect of the present embodiment, such latency may include a column address strobe (CAS) latency, which refers to a latency associated with accessing a column of data. Still yet, the latency may include a row address to column address latency (tRCD), which refers to a latency required between the row address strobe (RAS) and CAS. Even still, the latency may include a row precharge latency (tRP), which refers a latency required to terminate access to an open row, and open access to a next row. Further, the latency may include an activate to precharge latency (tRAS), which refers to a latency required to access a certain row of data between an activate operation and a precharge operation. In any case, the interface circuit **102** may be operable for simulating at least one memory circuit with a first latency that is longer (or shorter) than a second latency of at least one of the memory circuits **104A**, **104B**, **104N**. Different optional embodiments which employ various features of the present embodiment will be set forth hereinafter during reference to FIGS. **2A-2E**, and **3** et al.

[0033] In still another embodiment, a component may be operable to receive a signal from the system **106** and communicate the signal to at least one of the memory circuits **104A**, **104B**, **104N** after a delay. Again, the signal may refer to a control signal (e.g. an address signal; a signal associated with an activate operation, precharge operation, write operation, read operation; etc.), a data signal, a logical or physical signal, or any other signal for that matter. In various embodiments, such delay may be fixed or variable (e.g. a function of the current signal, the previous signal, etc.). In still other embodiments, the component may be operable to receive a signal from at least one of the memory circuits **104A**, **104B**, **104N** and communicate the signal to the system **106** after a delay.

[0034] As an option, the delay may include a cumulative delay associated with any one or more of the aforementioned signals. Even still, the delay may result in a time shift of the signal forward and/or back in time (with respect to other signals). Of course, such forward and backward time shift may or may not be equal in magnitude. In one embodiment, this time shifting may be accomplished by utilizing a plurality of delay functions which each apply a different delay to a different signal. In still additional embodiments, the aforementioned shifting may be coordinated among multiple signals such that different signals are subject to shifts with different relative directions/magnitudes, in an organized fashion.

[0035] Further, it should be noted that the aforementioned component may, but need not necessarily take the form of the interface circuit **102** of FIG. **1**. For example, the component may include a register, an AMB, a component positioned on at least one DIMM, a memory controller, etc.

Such register may, in various embodiments, include a Joint Electron Device Engineering Council (JEDEC) register, a JEDEC register including one or more functions set forth herein, a register with forwarding, storing, and/or buffering capabilities, etc. Different optional embodiments which employ various features of the present embodiment will be set forth hereinafter during reference to FIGS. **4-7**, and **9A-B** et al.

[0036] In a power-saving embodiment, at least one of a plurality of memory circuits **104A**, **104B**, **104N** may be identified that is not currently being accessed by the system **106**. In one embodiment, such identification may involve determining whether a page [i.e. any portion of any memory (s), etc.] is being accessed in at least one of the plurality of memory circuits **104A**, **104B**, **104N**. Of course, any other technique may be used that results in the identification of at least one of the memory circuits **104A**, **104B**, **104N** that is not being accessed.

[0037] In response to the identification of the at least one memory circuit **104A**, **104B**, **104N**, a power saving operation is initiated in association with the at least one memory circuit **104A**, **104B**, **104N**. In one optional embodiment, such power saving operation may involve a power down operation and, in particular, a precharge power down operation. Of course, however, it should be noted that any operation that results in at least some power savings may be employed in the context of the present embodiment.

[0038] Similar to one or more of the previous embodiments, the present functionality or a portion thereof may be carried out utilizing any desired component. For example, such component may, but need not necessarily take the form of the interface circuit **102** of FIG. **1**. In other embodiments, the component may include a register, an AMB, a component positioned on at least one DIMM, a memory controller, etc. One optional embodiment which employs various features of the present embodiment will be set forth hereinafter during reference to FIG. **10**.

[0039] In still yet another embodiment, a plurality of the aforementioned components may serve, in combination, to interface the memory circuits **104A**, **104B**, **104N** and the system **106**. In various embodiments, two, three, four, or more components may accomplish this. Also, the different components may be relatively configured in any desired manner. For example, the components may be configured in parallel, serially, or a combination thereof. In addition, any number of the components may be allocated to any number of the memory circuits **104A**, **104B**, **104N**.

[0040] Further, in the present embodiment, each of the plurality of components may be the same or different. Still yet, the components may share the same or similar interface tasks and/or perform different interface tasks. Such interface tasks may include, but are not limited to simulating one or more aspects of a memory circuit, performing a power savings/refresh operation, carrying out any one or more of the various functionalities set forth herein, and/or any other task relevant to the aforementioned interfacing. One optional embodiment which employs various features of the present embodiment will be set forth hereinafter during reference to FIG. **3**.

[0041] Additional illustrative information will now be set forth regarding various optional embodiments in which the foregoing techniques may or may not be implemented, per the desires of the user. For example, an embodiment is set forth for storing at least a portion of information received in

US 2008/0025137 A1

Jan. 31, 2008

4

association with a first operation for use in performing a second operation. See FIG. 2F. Further, a technique is provided for refreshing a plurality of memory circuits, in accordance with still yet another embodiment. See FIG. 11.

[0042] It should again be strongly noted that the following information is set forth for illustrative purposes and should not be construed as limiting in any manner. Any of the following features may be optionally incorporated with or without the exclusion of other features described.

[0043] FIGS. 2A-2E show various configurations of a buffered stack of DRAM circuits 206A-D with a buffer chip 202, in accordance with various embodiments. As an option, the various configurations to be described in the following embodiments may be implemented in the context of the architecture and/or environment of FIG. 1. Of course, however, they may also be carried out in any other desired environment (e.g. using other memory types, etc.). It should also be noted that the aforementioned definitions may apply during the present description.

[0044] As shown in each of such figures, the buffer chip 202 is placed electrically between an electronic host system 204 and a stack of DRAM circuits 206A-D. In the context of the present description, a stack may refer to any collection of memory circuits. Further, the buffer chip 202 may include any device capable of buffering a stack of circuits (e.g. DRAM circuits 206A-D, etc.). Specifically, the buffer chip 202 may be capable of buffering the stack of DRAM circuits 206A-D to electrically and/or logically resemble at least one larger capacity DRAM circuit to the host system 204. In this way, the stack of DRAM circuits 206A-D may appear as a smaller quantity of larger capacity DRAM circuits to the host system 204.

[0045] For example, the stack of DRAM circuits 206A-D may include eight 512 Mb DRAM circuits. Thus, the buffer chip 202 may buffer the stack of eight 512 Mb DRAM circuits to resemble a single 4 Gb DRAM circuit to a memory controller (not shown) of the associated host system 204. In another example, the buffer chip 202 may buffer the stack of eight 512 Mb DRAM circuits to resemble two 2 Gb DRAM circuits to a memory controller of an associated host system 204.

[0046] Further, the stack of DRAM circuits 206A-D may include any number of DRAM circuits. Just by way of example, a buffer chip 202 may be connected to 2, 4, 8 or more DRAM circuits 206A-D. Also, the DRAM circuits 206A-D may be arranged in a single stack, as shown in FIGS. 2A-2D.

[0047] The DRAM circuits 206A-D may be arranged on a single side of the buffer chip 202, as shown in FIGS. 2A-2D. Of course, however, the DRAM circuits 206A-D may be located on both sides of the buffer chip 202 shown in FIG. 2E. Thus, for example, a buffer chip 202 may be connected to 16 DRAM circuits with 8 DRAM circuits on either side of the buffer chip 202, where the 8 DRAM circuits on each side of the buffer chip 202 are arranged in two stacks of four DRAM circuits.

[0048] The buffer chip 202 may optionally be a part of the stack of DRAM circuits 206A-D. Of course, however, the buffer chip 202 may also be separate from the stack of DRAM circuits 206A-D. In addition, the buffer chip 202 may be physically located anywhere in the stack of DRAM circuits 206A-D, where such buffer chip 202 electrically sits between the electronic host system 204 and the stack of DRAM circuits 206A-D.

[0049] In one embodiment, a memory bus (not shown) may connect to the buffer chip 202, and the buffer chip 202 may connect to each of the DRAM circuits 206A-D in the stack. As shown in FIGS. 2A-D, the buffer chip 202 may be located at the bottom of the stack of DRAM circuits 206A-D (e.g. the bottom-most device in the stack). As another option, and as shown in FIG. 2E, the buffer chip 202 may be located in the middle of the stack of DRAM circuits 206A-D. As still yet another option, the buffer chip 202 may be located at the top of the stack of DRAM circuits 206A-D (e.g. the top-most device in the stack). Of course, however, the buffer chip 202 may be located anywhere between the two extremities of the stack of DRAM circuits 206A-D.

[0050] The electrical connections between the buffer chip 202 and the stack of DRAM circuits 206A-D may be configured in any desired manner. In one optional embodiment; address, control (e.g. command, etc.), and clock signals may be common to all DRAM circuits 206A-D in the stack (e.g. using one common bus). As another option, there may be multiple address, control and clock busses. As yet another option, there may be individual address, control and clock busses to each DRAM circuit 206A-D. Similarly, data signals may be wired as one common bus, several busses or as an individual bus to each DRAM circuit 206A-D. Of course, it should be noted that any combinations of such configurations may also be utilized.

[0051] For example, as shown in FIG. 2A, the stack of DRAM circuits 206A-D may have one common address, control and clock bus 208 with individual data busses 210. In another example, as shown in FIG. 2B, the stack of DRAM circuits 206A-D may have two address, control and clock busses 208 along with two data busses 210. In still yet another example, as shown in FIG. 2C, the stack of DRAM circuits 206A-D may have one address, control and clock bus 208 together with two data busses 210. In addition, as shown in FIG. 2D, the stack of DRAM circuits 206A-D may have one common address, control and clock bus 208 and one common data bus 210. It should be noted that any other permutations and combinations of such address, control, clock and data busses may be utilized.

[0052] These configurations may therefore allow for the host system 204 to only be in contact with a load of the buffer chip 202 on the memory bus. In this way, any electrical loading problems (e.g. bad signal integrity, improper signal timing, etc.) associated with the stacked DRAM circuits 206A-D may (but not necessarily) be prevented, in the context of various optional embodiments.

[0053] FIG. 2F illustrates a method 280 for storing at least a portion of information received in association with a first operation for use in performing a second operation, in accordance with still yet another embodiment. As an option, the method 280 may be implemented in the context of the architecture and/or environment of any one or more of FIGS. 1-2E. For example, the method 280 may be carried out by the interface circuit 102 of FIG. 1. Of course, however, the method 280 may be carried out in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0054] In operation 282, first information is received in association with a first operation to be performed on at least one of a plurality of memory circuits (e.g. see the memory circuits 104A, 104B, 104N of FIG. 1, etc.). In various embodiments, such first information may or may not be received coincidentally with the first operation, as long as it is

US 2008/0025137 A1

Jan. 31, 2008

5

associated in some capacity. Further, the first operation may, in one embodiment, include a row operation. In such embodiment, the first information may include address information (e.g. a set of address bits, etc.).

[0055] For reasons that will soon become apparent, at least a portion of the first information is stored. Note operation 284. Still yet, in operation 286, second information is received in association with a second operation. Similar to the first information, the second information may or may not be received coincidentally with the second operation, and may include address information. Such second operation, however, may, in one embodiment, include a column operation. [0056] To this end, the second operation may be performed utilizing the stored portion of the first information in addition to the second information. See operation 288. More illustrative information will now be set forth regarding various optional features with which the foregoing method 280 may or may not be implemented, per the desires of the user. Specifically, an example will be set forth illustrating the manner in which the method 280 may be employed for accommodating a buffer chip that is simulating at least one aspect of a plurality of memory circuits.

[0057] In particular, the present example of the method 280 of FIG. 2F will be set forth in the context of the various components (e.g. buffer chip 202, etc.) shown in the embodiments of FIGS. 2A-2E. It should be noted that, since the buffered stack of DRAM circuits 206A-D may appear to the memory controller of the host system 204 as one or more larger capacity DRAM circuits, the buffer chip 202 may receive more address bits from the memory controller than are required by the DRAM circuits 206A-D in the stack. These extra address bits may be decoded by the buffer chip 202 to individually select the DRAM circuits 206A-D in the stack, utilizing separate chip select signals to each of the DRAM circuits 206A-D in the stack.

[0058] For example, a stack of four $\times 4$ 1 Gb DRAM circuits 206A-D behind a buffer chip 202 may appear as a single $\times 4$ 4 Gb DRAM circuit to the memory controller. Thus, the memory controller may provide sixteen row address bits and three bank address bits during a row (e.g. activate) operation, and provide eleven column address bits and three bank address bits during a column (e.g. read or write) operation. However, the individual DRAM circuits 206A-D in the stack may require only fourteen row address bits and three bank address bits for a row operation, and eleven column address bits and three bank address bits during a column operation.

[0059] As a result, during a row operation in the above example, the buffer chip 202 may receive two address bits more than are needed by each DRAM circuit 206A-D in the stack. The buffer chip 202 may therefore use the two extra address bits from the memory controller to select one of the four DRAM circuits 206A-D in the stack. In addition, the buffer chip 202 may receive the same number of address bits from the memory controller during a column operation as are needed by each DRAM circuit 206A-D in the stack.

[0060] Thus, in order to select the correct DRAM circuit 206A-D in the stack during a column operation, the buffer chip 202 may be designed to store the two extra address bits provided during a row operation and use the two stored address bits to select the correct DRAM circuit 206A-D during the column operation. The mapping between a system address (e.g. address from the memory controller, including the chip select signal(s)) and a device address (e.g.

the address, including the chip select signals, presented to the DRAM circuits 206A-D in the stack) may be performed by the buffer chip 202 in various manners.

[0061] In one embodiment, a lower order system row address and bank address bits may be mapped directly to the device row address and bank address inputs. In addition, the most significant row address bit(s) and, optionally, the most significant bank address bit(s), may be decoded to generate the chip select signals for the DRAM circuits 206A-D in the stack during a row operation. The address bits used to generate the chip select signals during the row operation may also be stored in an internal lookup table by the buffer chip 22 for one or more clock cycles. During a column operation, the system column address and bank address bits may be mapped directly to the device column address and bank address inputs, while the stored address bits may be decoded to generate the chip select signals.

[0062] For example, addresses may be mapped between four 512 Mb DRAM circuits 206A-D that simulate a single 2 Gb DRAM circuits utilizing the buffer chip 202. There may be 15 row address bits from the system 204, such that row address bits 0 through 13 are mapped directly to the DRAM circuits 206A-D. There may also be 3 bank address bits from the system 204, such that bank address bits 0 through 1 are mapped directly to the DRAM circuits 206A-D.

[0063] During a row operation, the bank address bit 2 and the row address bit 14 may be decoded to generate the 4 chip select signals for each of the four DRAM circuits 206A-D. Row address bit 14 may be stored during the row operation using the bank address as the index. In addition, during the column operation, the stored row address bit 14 may again be used with bank address bit 2 to form the four DRAM chip select signals.

[0064] As another example, addresses may be mapped between four 1 Gb DRAM circuits 206A-D that simulate a single 4 Gb DRAM circuits utilizing the buffer chip 202. There may be 16 row address bits from the system 204, such that row address bits 0 through 14 are mapped directly to the DRAM circuits 206A-D. There may also be 3 bank address bits from the system 204, such that bank address bits 0 through 3 are mapped directly to the DRAM circuits 206A-D.

[0065] During a row operation, row address bits 14 and 15 may be decoded to generate the 4 chip select signals for each of the four DRAM circuits 206A-D. Row address bits 14 and 15 may also be stored during the row operation using the bank address as the index. During the column operation, the stored row address bits 14 and 15 may again be used to form the four DRAM chip select signals.

[0066] In various embodiments, this mapping technique may optionally be used to ensure that there are no unnecessary combinational logic circuits in the critical timing path between the address input pins and address output pins of the buffer chip 202. Such combinational logic circuits may instead be used to generate the individual chip select signals. This may therefore allow the capacitive loading on the address outputs of the buffer chip 202 to be much higher than the loading on the individual chip select signal outputs of the buffer chip 202.

[0067] In another embodiment, the address mapping may be performed by the buffer chip 202 using some of the bank address signals from the memory controller to generate the individual chip select signals. The buffer chip 202 may store

US 2008/0025137 A1

Jan. 31, 2008

6

the higher order row address bits during a row operation using the bank address as the index, and then may use the stored address bits as part of the DRAM circuit bank address during a column operation. This address mapping technique may require an optional lookup table to be positioned in the critical timing path between the address inputs from the memory controller and the address outputs, to the DRAM circuits 206A-D in the stack.

[0068] For example, addresses may be mapped between four 512 Mb DRAM circuits 206A-D that simulate a single 2 Gb DRAM utilizing the buffer chip 202. There may be 15 row address bits from the system 204, where row address bits 0 through 13 are mapped directly to the DRAM circuits 206A-D. There may also be 3 bank address bits from the system 204, such that bank address bit 0 is used as a DRAM circuit bank address bit for the DRAM circuits 206A-D.

[0069] In addition, row address bit 14 may be used as an additional DRAM circuit bank address bit. During a row operation, the bank address bits 1 and 2 from the system may be decoded to generate the 4 chip select signals for each of the four DRAM circuits 206A-D. Further, row address bit 14 may be stored during the row operation. During the column operation, the stored row address bit 14 may again be used along with the bank address bit 0 from the system to form the DRAM circuit bank address.

[0070] In both of the above described address mapping techniques, the column address from the memory controller may be mapped directly as the column address to the DRAM circuits 206A-D in the stack. Specifically, this direct mapping may be performed since each of the DRAM circuits 206A-D in the stack, even if of the same width but different capacities (e.g. from 512 Mb to 4 Gb), may have the same page sizes. In an optional embodiment, address A[10] may be used by the memory controller to enable or disable auto-precharge during a column operation. Therefore, the buffer chip 202 may forward A[10] from the memory controller to the DRAM circuits 206A-D in the stack without any modifications during a column operation.

[0071] In various embodiments, it may be desirable to determine whether the simulated DRAM circuit behaves according to a desired DRAM standard or other design specification. A behavior of many DRAM circuits is specified by the JEDEC standards and it may be desirable, in some embodiments, to exactly simulate a particular JEDEC standard DRAM. The JEDEC standard defines control signals that a DRAM circuit must accept and the behavior of the DRAM circuit as a result of such control signals. For example, the JEDEC specification for a DDR2 DRAM is known as JESD79-2B.

[0072] If it is desired, for example, to determine whether a JEDEC standard is met, the following algorithm may be used. Such algorithm checks, using a set of software verification tools for formal verification of logic, that protocol behavior of the simulated DRAM circuit is the same as a desired standard or other design specification. This formal verification is quite feasible because the DRAM protocol described in a DRAM standard is typically limited to a few control signals (e.g. approximately 15 control signals in the case of the JEDEC DDR2 specification, for example).

[0073] Examples of the aforementioned software verification tools include MAGELLAN supplied by SYNOPSIS, or other software verification tools, such as INCISIVE supplied by CADENCE, verification tools supplied by JASPER, VERIX supplied by REAL INTENT, 0-IN supplied by

MENTOR CORPORATION, and others. These software verification tools use written assertions that correspond to the rules established by the DRAM protocol and specification. These written assertions are further included in the code that forms the logic description for the buffer chip. By writing assertions that correspond to the desired behavior of the simulated DRAM circuit, a proof may be constructed that determines whether the desired design requirements are met. In this way, one may test various embodiments for compliance with a standard, multiple standards, or other design specification.

[0074] For instance, an assertion may be written that no two DRAM control signals are allowed to be issued to an address, control and clock bus at the same time. Although one may know which of the various buffer chip/DRAM stack configurations and address mappings that have been described herein are suitable, the aforementioned algorithm may allow a designer to prove that the simulated DRAM circuit exactly meets the required standard or other design specification. If, for example, an address mapping that uses a common bus for data and a common bus for address results in a control and clock bus that does not meet a required specification, alternative designs for buffer chips with other bus arrangements or alternative designs for the interconnect between the buffer chips may be used and tested for compliance with the desired standard or other design specification.

[0075] FIG. 3 shows a high capacity DIMM 300 using buffered stacks of DRAM circuits 302, in accordance with still yet another embodiment. As an option, the high capacity DIMM 300 may be implemented in the context of the architecture and environment of FIGS. 1 and/or 2A-F. Of course, however, the high capacity DIMM 300 may be used in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0076] As shown, a high capacity DIMM 300 may be created utilizing buffered stacks of DRAM circuits 302. Thus, a DIMM 300 may utilize a plurality of buffered stacks of DRAM circuits 302 instead of individual DRAM circuits, thus increasing the capacity of the DIMM. In addition, the DIMM 300 may include a register 304 for address and operation control of each of the buffered stacks of DRAM circuits 302. It should be noted that any desired number of buffered stacks of DRAM circuits 302 may be utilized in conjunction with the DIMM 300. Therefore, the configuration of the DIMM 300, as shown, should not be construed as limiting in any way.

[0077] In an additional unillustrated embodiment, the register 304 may be substituted with an AMB (not shown), in the context of an FB-DIMM.

[0078] FIG. 4 shows a timing design 400 of a buffer chip that makes a buffered stack of DRAM circuits mimic longer CAS latency DRAM to a memory controller, in accordance with another embodiment. As an option, the design of the buffer chip may be implemented in the context of the architecture and environment of FIGS. 1-3. Of course, however, the design of the buffer chip may be used in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0079] In use, any delay through a buffer chip (e.g. see the buffer chip 202 of FIGS. 2A-E, etc.) may be made transparent to a memory controller of a host system (e.g. see the

US 2008/0025137 A1

Jan. 31, 2008

7

host system **204** of FIGS. 2A-E, etc.) utilizing the buffer chip. In particular, the buffer chip may buffer a stack of DRAM circuits such that the buffered stack of DRAM circuits appears as at least one larger capacity DRAM circuit with higher CAS latency.

[0080] Such delay may be a result of the buffer chip being located electrically between the memory bus of the host system and the stacked DRAM circuits, since most or all of the signals that connect the memory bus to the DRAM circuits pass through the buffer chip. A finite amount of time may therefore be needed for these signals to traverse through the buffer chip. With the exception of register chips and advanced memory buffers (AMB), industry standard protocols for memory [e.g. (DDR SDRAM), DDR2 SDRAM, etc.] may not comprehend the buffer chip that sits between the memory bus and the DRAM. Industry standard protocols for memory [e.g. (DDR SDRAM), DDR2 SDRAM, etc.] narrowly define the properties of chips that sit between host and memory circuits. Such industry standard protocols define the properties of a register chip and AMB but not the properties of the buffer chip **202**, etc. Thus, the signal delay through the buffer chip may violate the specifications of industry standard protocols.

[0081] In one embodiment, the buffer chip may provide a one-half clock cycle delay between the buffer chip receiving address and control signals from the memory controller (or optionally from a register chip, an AMB, etc.) and the address and control signals being valid at the inputs of the stacked DRAM circuits. Similarly, the data signals may also have a one-half clock cycle delay in traversing the buffer chip, either from the memory controller to the DRAM circuits or from the DRAM circuits to the memory controller. Of course, the one-half clock cycle delay set forth above is set forth for illustrative purposes only and thus should not be construed as limiting in any manner whatsoever. For example, other embodiments are contemplated where a one clock cycle delay, a multiple clock cycle delay (or fraction thereof), and/or any other delay amount is incorporated, for that matter. As mentioned earlier, in other embodiments, the aforementioned delay may be coordinated among multiple signals such that different signals are subject to time-shifting with different relative directions/magnitudes, in an organized fashion.

[0082] As shown in FIG. 4, the cumulative delay through the buffer chip (e.g. the sum of a first delay **402** of the address and control signals through the buffer chip and a second delay **404** of the data signals through the buffer chip) is j clock cycles. Thus, the buffer chip may make the buffered stack appear to the memory controller as one or more larger DRAM circuits with a CAS latency **408** of $i+j$ clocks, where i is the native CAS latency of the DRAM circuits.

[0083] In one example, if the DRAM circuits in the stack have a native CAS latency of 4 and the address and control signals along with the data signals experience a one-half clock cycle delay through the buffer chip, then the buffer chip may make the buffered stack appear to the memory controller as one or more larger DRAM circuits with a CAS latency of 5 (i.e. $4+1$). In another example, if the address and control signals along with the data signals experience a 1 clock cycle delay through the buffer chip, then the buffer chip may make the buffered stack appear as one or more larger DRAM circuits with a CAS latency of 6 (i.e. $4+2$).

[0084] FIG. 5 shows the write data timing **500** expected by a DRAM circuit in a buffered stack, in accordance with yet another embodiment. As an option, the write data timing **500** may be implemented in the context of the architecture and environment of FIGS. 1-4. Of course, however, the write data timing **500** may be carried out in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0085] Designing a buffer chip (e.g. see the buffer chip **202** of FIGS. 2A-E, etc.) so that a buffered stack appears as at least one larger capacity DRAM circuit with higher CAS latency may, in some embodiments, create a problem with the timing of write operations. For example, with respect to a buffered stack of DDR2 SDRAM circuits with a CAS latency of 4 that appear as a single larger DDR2 SDRAM with a CAS latency of 6 to the memory controller, the DDR2 SDRAM protocol may specify that the write CAS latency is one less than the read CAS latency. Therefore, since the buffered stack appears as a DDR2 SDRAM with a read CAS latency of 6, the memory controller may use a write CAS latency of 5 (see **502**) when scheduling a write operation to the buffered stack.

[0086] However, since the native read CAS latency of the DRAM circuits is 4, the DRAM circuits may require a write CAS latency of 3 (see **504**). As a result, the write data from the memory controller may arrive at the buffer chip later than when the DRAM circuits require the data. Thus, the buffer chip may delay such write operations to alleviate any of such timing problems. Such delay in write operations will be described in more detail with respect to FIG. 6 below.

[0087] FIG. 6 shows write operations **600** delayed by a buffer chip, in accordance with still yet another embodiment. As an option, the write operations **600** may be implemented in the context of the architecture and environment of FIGS. 1-5. Of course, however, the write operations **600** may be used in any desired environment. Again, it should also be noted that the aforementioned definitions may apply during the present description.

[0088] In order to be compliant with the protocol utilized by the DRAM circuits in the stack, a buffer chip (e.g. see the buffer chip **202** of FIGS. 2A-E, etc.) may provide an additional delay, over and beyond the delay of the address and control signals through the buffer chip, between receiving the write operation and address from the memory controller (and/or optionally from a register and/or AMB, etc.), and sending it to the DRAM circuits in the stack. The additional delay may be equal to j clocks, where j is the cumulative delay of the address and control signals through the buffer chip and the delay of the data signals through the buffer chip. As another option, the write address and operation may be delayed by a register chip on a DIMM, by a AMB, or by the memory controller.

[0089] FIG. 7 shows early write data **700** from an AMB, in accordance with another embodiment. As an option, the early write data **700** may be implemented in the context of the architecture and environment of FIGS. 1-5. Of course, however, the early write data **700** may be used in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0090] As shown, an AMB on an FB-DIMM may be designed to send write data earlier to buffered stacks instead of delaying the write address and operation, as described in reference to FIG. 6. Specifically, an early write latency **702** may be utilized to send the write data to the buffered stack.

US 2008/0025137 A1

Jan. 31, 2008

8

Thus, correct timing of the write operation at the inputs of the DRAM circuits in the stack may be ensured.

[0091] For example, a buffer chip (e.g. see the buffer chip 202 of FIGS. 2A-E, etc.) may have a cumulative latency of 2, in which case, the AMB may send the write data 2 clock cycles earlier to the buffered stack. It should be noted that this scheme may not be possible in the case of registered DIMMs since the memory controller sends the write data directly to the buffered stacks. As an option, a memory controller may be designed to send write data earlier so that write operations have the correct timing at the input of the DRAM circuits in the stack without requiring the buffer chip to delay the write address and operation.

[0092] FIG. 8 shows address bus conflicts 800 caused by delayed write operations, in accordance with yet another embodiment. As mentioned earlier, the delaying of the write addresses and operations may be performed by a buffer chip, or optionally a register, AMB, etc., in a manner that is completely transparent to the memory controller of a host system. However, since the memory controller is unaware of this delay, it may schedule subsequent operations, such as for example activate or precharge operations, which may collide with the delayed writes on the address bus from the buffer chip to the DRAM circuits in the stack. As shown, an activate operation 802 may interfere with a write operation 804 that has been delayed. Thus, a delay of activate operations may be employed, as will be described in further detail with respect to FIG. 9.

[0093] FIGS. 9A-B show variable delays 900 and 950 of operations through a buffer chip, in accordance with another embodiment. As an option, the variable delays 900 and 950 may be implemented in the context of the architecture and environment of FIGS. 1-8. Of course, however, the variable delays 900 and 950 may be carried out in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0094] In order to prevent conflicts on an address bus between the buffer chip and its associated stack(s), either the write operation or the precharge/activate operation may be delayed. As shown, a buffer chip (e.g. see the buffer chip 202 of FIGS. 2A-E, etc.) may delay the precharge/activate operations 952A-C/902A-C. In particular, the buffer chip may make the buffered stack appear as one or more larger capacity DRAM circuits that have longer tRCD (RAS to CAS delay) and tRP (i.e. precharge time) parameters.

[0095] For example, if the cumulative latency through a buffer chip is 2 clock cycles while the native read CAS latency of the DRAM circuits is 4 clock cycles, then in order to hide the delay of the address/control signals and the data signals through the buffer chip, the buffered stack may appear as one or more larger capacity DRAM circuits with a read CAS latency of 6 clock cycles to the memory controller. In addition, if the tRCD and tRP of the DRAM circuits is 4 clock cycles each, the buffered stack may appear as one or more larger capacity DRAM circuits with tRCD of 6 clock cycles and tRP of 6 clock cycles in order to allow a buffer chip (e.g., see the buffer chip 202 of FIGS. 2A-E, etc.) to delay the activate and precharge operations in a manner that is transparent to the memory controller. Specifically, a buffered stack that uses 4-4-4 DRAM circuits (i.e. CAS latency=4, tRCD=4, tRP=4) may appear as one or at least one larger capacity DRAM circuits with 6-6-6 timing (i.e. CAS latency=6, tRCD=6, tRP=6).

[0096] Since the buffered stack appears to the memory controller as having a tRCD of 6 clock cycles, the memory controller may schedule a column operation to a bank 6 clock cycles after an activate (e.g. row) operation to the same bank. However, the DRAM circuits in the stack may actually have a tRCD of 4 clock cycles. Thus, the buffer chip may have the ability to delay the activate operation by up to 2 clock cycles in order to avoid any conflicts on the address bus between the buffer chip and the DRAM circuits in the stack while still ensuring correct read and write timing on the channel between the memory controller and the buffered stack.

[0097] As shown, the buffer chip may issue the activate operation to the DRAM circuits one, two, or three clock cycles after it receives the activate operation from the memory controller, register, or AMB. The actual delay of the activate operation through the buffer chip may depend on the presence or absence of other DRAM operations that may conflict with the activate operation, and may optionally change from one activate operation to another.

[0098] Similarly, since the buffered stack may appear to the memory controller as at least one larger capacity DRAM circuit with a tRP of 6 clock cycles, the memory controller may schedule a subsequent activate (e.g. row) operation to a bank a minimum of 6 clock cycles after issuing a precharge operation to that bank. However, since the DRAM circuits in the stack actually have a tRP of 4 clock cycles, the buffer chip may have the ability to delay issuing the precharge operation to the DRAM circuits in the stack by up to 2 clock cycles in order to avoid any conflicts on the address bus between the buffer chip and the DRAM circuits in the stack. In addition, even if there are no conflicts on the address bus, the buffer chip may still delay issuing a precharge operation in order to satisfy the tRAS requirement of the DRAM circuits.

[0099] In particular, if the activate operation to a bank was delayed to avoid an address bus conflict, then the precharge operation to the same bank may be delayed by the buffer chip to satisfy the tRAS requirement of the DRAM circuits. The buffer chip may issue the precharge operation to the DRAM circuits one, two, or three clock cycles after it receives the precharge operation from the memory controller, register, or AMB. The actual delay of the precharge operation through the buffer chip may depend on the presence or absence of address bus conflicts or tRAS violations, and may change from one precharge operation to another.

[0100] FIG. 10 shows a buffered stack 1000 of four 512 Mb DRAM circuits mapped to a single 2 Gb DRAM circuit, in accordance with yet another embodiment. As an option, the buffered stack 1000 may be implemented in the context of the architecture and environment of FIGS. 1-9. Of course, however, the buffered stack 1000 may be carried out in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0101] The multiple DRAM circuits 1002A-D buffered in the stack by the buffer chip 1004 may appear as at least one larger capacity DRAM circuit to the memory controller. However, the combined power dissipation of such DRAM circuits 1002A-D may be much higher than the power dissipation of a monolithic DRAM of the same capacity. For example, the buffered stack may consist of four 512 Mb DDR2 SDRAM circuits that appear to the memory controller as a single 2 Gb DDR2 SDRAM circuit.

US 2008/0025137 A1

Jan. 31, 2008

9

[0102] The power dissipation of all four DRAM circuits 1002A-D in the stack may be much higher than the power dissipation of a monolithic 2 Gb DDR2 SDRAM. As a result, a DIMM containing multiple buffered stacks may dissipate much more power than a standard DIMM built using monolithic DRAM circuits. This increased power dissipation may limit the widespread adoption of DIMMs that use buffered stacks.

[0103] Thus, a power management technique that reduces the power dissipation of DIMMs that contain buffered stacks of DRAM circuits may be utilized. Specifically, the DRAM circuits 1002A-D may be opportunistically placed in a precharge power down mode using the clock enable (CKE) pin of the DRAM circuits 1002A-D. For example, a single rank registered DIMM (R-DIMM) may contain a plurality of buffered stacks of DRAM circuits 1002A-D, where each stack consists of four $\times 4$ 512 Mb DDR2 SDRAM circuits 1002A-D and appears as a single $\times 4$ 2 Gb DDR2 SDRAM circuit to the memory controller. A 2 Gb DDR2 SDRAM may generally have eight banks as specified by JEDEC. Therefore, the buffer chip 1004 may map each 512 Mb DRAM circuit in the stack to two banks of the equivalent 2 Gb DRAM, as shown.

[0104] The memory controller of the host system may open and close pages in the banks of the DRAM circuits 1002A-D based on the memory requests it receives from the rest of the system. In various embodiments, no more than one page may be able to be open in a bank at any given time. For example, with respect to FIG. 10, since each DRAM circuit 1002A-D in the stack is mapped to two banks of the equivalent larger DRAM, at any given time a DRAM circuit 1002A-D may have two open pages, one open page, or no open pages. When a DRAM circuit 1002A-D has no open pages, the power management scheme may place that DRAM circuit 1002A-D in the precharge power down mode by de-asserting its CKE input.

[0105] The CKE inputs of the DRAM circuits 1002A-D in a stack may be controlled by the buffer chip 1004, by a chip on an R-DIMM, by an AMB on a FB-DIMM, or by the memory controller in order to implement the power management scheme described hereinabove. In one embodiment, this power management scheme may be particularly efficient when the memory controller implements a closed page policy.

[0106] Another optional power management scheme may include mapping a plurality of DRAM circuits to a single bank of the larger capacity DRAM seen by the memory controller. For example, a buffered stack of sixteen $\times 4$ 256 Mb DDR2 SDRAM circuits may appear to the memory controller as a single $\times 4$ 4 Gb DDR2 SDRAM circuit. Since a 4 Gb DDR2 SDRAM circuit is specified by JEDEC to have eight banks, each bank of the 4 Gb DDR2 SDRAM circuit may be 512 Mb. Thus, two of the 256 Mb DDR2 SDRAM circuits may be mapped by the buffer chip 1004 to a single bank of the equivalent 4 Gb DDR2 SDRAM circuit seen by the memory controller.

[0107] In this way, bank 0 of the 4 Gb DDR2 SDRAM circuit may be mapped by the buffer chip to two 256 Mb DDR2 SDRAM circuits (e.g. DRAM A and DRAM B) in the stack. However, since only one page can be open in a bank at any given time, only one of DRAM A or DRAM B may be in the active state at any given time. If the memory controller opens a page in DRAM A, then DRAM B may be placed in the precharge power down mode by de-asserting

its CKE input. As another option, if the memory controller opens a page in DRAM B, DRAM A may be placed in the precharge power down mode by de-asserting its CKE input. This technique may ensure that if p DRAM circuits are mapped to a bank of the larger capacity DRAM circuit seen by the memory controller, then p-1 of the p DRAM circuits may continuously (e.g. always, etc.) be subjected to a power saving operation. The power saving operation may, for example, comprise operating in precharge power down mode except when refresh is required. Of course, power-savings may also occur in other embodiments without such continuity.

[0108] FIG. 11 illustrates a method 1100 for refreshing a plurality of memory circuits, in accordance with still yet another embodiment. As an option, the method 1100 may be implemented in the context of the architecture and environment of any one or more of FIGS. 1-10. For example, the method 1100 may be carried out by the interface circuit 102 of FIG. 1. Of course, however, the method 1100 may be carried out in any desired environment. It should also be noted that the aforementioned definitions may apply during the present description.

[0109] As shown, a refresh control signal is received in operation 1102. In one optional embodiment, such refresh control signal may, for example, be received from a memory controller, where such memory controller intends to refresh a simulated memory circuit(s).

[0110] In response to the receipt of such refresh control signal, a plurality of refresh control signals are sent to a plurality of the memory circuits (e.g. see the memory circuits 104A, 104B, 104N of FIG. 1, etc.), at different times. See operation 1104. Such refresh control signals may or may not each include the refresh control signal of operation 1102 or an instantiation/copy thereof. Of course, in other embodiments, the refresh control signals may each include refresh control signals that are different in at least one aspect (e.g. format, content, etc.).

[0111] During use of still additional embodiments, at least one first refresh control signal may be sent to a first subset (e.g. of one or more) of the memory circuits at a first time and at least one second refresh control signal may be sent to a second subset (e.g. of one or more) of the memory circuits at a second time. Thus, in some embodiments, a single refresh control signal may be sent to a plurality of the memory circuits (e.g. a group of memory circuits, etc.). Further, a plurality of the refresh control signals may be sent to a plurality of the memory circuits. To this end, refresh control signals may be sent individually or to groups of memory circuits, as desired.

[0112] Thus, in still yet additional embodiments, the refresh control signals may be sent after a delay in accordance with a particular timing. In one embodiment, for example, the timing in which the refresh control signals are sent to the memory circuits may be selected to minimize a current draw. This may be accomplished in various embodiments by staggering a plurality of refresh control signals. In still other embodiments, the timing in which the refresh control signals are sent to the memory circuits may be selected to comply with a tRFC parameter associated with each of the memory circuits.

[0113] To this end, in the context of an example involving a plurality of DRAM circuits (e.g. see the embodiments of FIGS. 1-2E, etc.), DRAM circuits of any desired size may receive periodic refresh operations to maintain the integrity

US 2008/0025137 A1

Jan. 31, 2008

10

of data therein. A memory controller may initiate refresh operations by issuing refresh control signals to the DRAM circuits with sufficient frequency to prevent any loss of data in the DRAM circuits. After a refresh control signal is issued to a DRAM circuit, a minimum time (e.g. denoted by tRFC) may be required to elapse before another control signal may be issued to that DRAM circuit. The tRFC parameter may therefore increase as the size of the DRAM circuit increases.

[0114] When the buffer chip receives a refresh control signal from the memory controller, it may refresh the smaller DRAM circuits within the span of time specified by the tRFC associated with the emulated DRAM circuit. Since the tRFC of the emulated DRAM circuits is larger than that of the smaller DRAM circuits, it may not be necessary to issue refresh control signals to all of the smaller DRAM circuits simultaneously. Refresh control signals may be issued separately to individual DRAM circuits or may be issued to groups of DRAM circuits, provided that the tRFC requirement of the smaller DRAM circuits is satisfied by the time the tRFC of the emulated DRAM circuits has elapsed. In use, the refreshes may be spaced to minimize the peak current draw of the combination buffer chip and DRAM circuit set during a refresh operation.

[0115] While various embodiments have been described above, it should be understood that they have been presented by way of example only, and not limitation. For example, any of the network elements may employ any of the desired functionality set forth hereinabove. Thus, the breadth and scope of a preferred embodiment should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

What is claimed is:

1. A sub-system, comprising:
an interface circuit in communication with a plurality of memory circuits and a system, the interface circuit operable to interface the memory circuits and the system for simulating at least one memory circuit with at least one aspect that is different from at least one aspect of at least one of the plurality of memory circuits;
wherein the at least one aspect is selected from the group consisting of a signal, a capacity, a timing, and a logical interface.
2. The sub-system of claim 1, wherein the at least one aspect includes the signal.
3. The sub-system of claim 2, wherein the signal is selected from the group consisting of an address signal, a control signal, and a data signal.
4. The sub-system of claim 1, wherein the at least one aspect includes the capacity.
5. The sub-system of claim 4, wherein the capacity includes a memory capacity.
6. The sub-system of claim 5, wherein memory capacity relates to a number of the memory circuits.
7. The sub-system of claim 5, wherein the interface circuit is operable to interface the memory circuits and the system

for simulating at least one memory circuit with a first memory capacity that is greater than a second memory capacity of at least one of the plurality of memory circuits.

8. The sub-system of claim 1, wherein the at least one aspect includes the timing.

9. The sub-system of claim 8, wherein the timing relates to a latency.

10. The sub-system of claim 9, wherein the latency includes a column address strobe (CAS) latency.

11. The sub-system of claim 10, wherein the interface circuit is operable to interface the memory circuits and the system for simulating at least one memory circuit with a first CAS latency that is longer than a second CAS latency of at least one of the plurality of memory circuits.

12. The sub-system of claim 9, wherein the latency includes a row address strobe (RAS) to CAS latency (tRCD).

13. The sub-system of claim 12, wherein the interface circuit is operable to interface the memory circuits and the system for simulating at least one memory circuit with a first tRCD that is longer than a second tRCD of at least one of the plurality of memory circuits.

14. The sub-system of claim 9, wherein the latency includes a row precharge latency (tRP).

15. The sub-system of claim 14, wherein the interface circuit is operable to interface the memory circuits and the system for simulating at least one memory circuit with a first tRP that is longer than a second tRP of at least one of the plurality of memory circuits.

16. The sub-system of claim 9, wherein the latency includes an activate to precharge latency (tRAS).

17. The sub-system of claim 1, wherein the at least one aspect includes the logical interface.

18. The sub-system of claim 1, wherein the memory circuits each include dynamic random access memory (DRAM).

19. A method, comprising:

interfacing a plurality of memory circuits and a system;
and

simulating at least one memory circuit with at least one aspect that is different from at least one aspect of at least one of the plurality of memory circuits;

wherein the at least one aspect is selected from the group consisting of a signal, a capacity, a timing, and a logical interface.

20. A sub-system, comprising:

an interface circuit in communication with a plurality of memory circuits and a system, the interface circuit operable to interface the memory circuits and the system for simulating at least one memory circuit with at least one aspect that is different from at least one aspect of at least one of the plurality of memory circuits;

wherein the interface circuit interfaces a majority of address or control signals of the memory circuits.

* * * * *